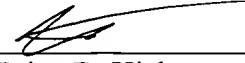




including extension of time fees, to Deposit Account 50-1302, and please credit any excess fees to such deposit account.

Respectfully submitted,

HICKMAN PALERMO TRUONG & BECKER LLP



Brian D. Hickman
Registration No. 35,894

Dated: April 10, 2001

1600 Willow Street
San Jose, California 95125-5106
Tel: (408) 414-1080 ext. 201
Fax: (408) 414-1076



Attorney Docket No. 50277-1646

MARKED-UP VERSION OF AMENDED CLAIMS

1 20. (Unchanged) A method of parallelizing an operation, the method comprising the steps of:
2 dividing the operation into a set of work partitions;
3 assigning work partitions from said set of work partitions to a plurality of entities,
4 wherein at least one entity of said plurality of entities is assigned a plurality of
5 work partitions from said set of work partitions; and
6 said plurality of entities operating in parallel on work partitions assigned to them to
7 perform said operation.

1 21. (Unchanged) The method of Claim 20 wherein:
2 the method further includes the step of receiving a query that requires the operation; and
3 the step of dividing the operation is performed in response to processing the query.

1 22. (Unchanged) The method of Claim 20 wherein the step of assigning work partitions
2 includes:
3 assigning said at least one entity a first work partition from said set of work partitions;
4 and
5 after said at least one entity has completed operating on said first work partition,
6 assigning said at least one entity a second work partition from said set of work
7 partitions.

1 23. (Unchanged) The method of Claim 20 wherein said plurality of entities are a plurality of
2 processes.

1 24. (Unchanged) The method of Claim 23 wherein said plurality of processes reside within a
2 single database system.

1 25. (Unchanged) The method of Claim 20 wherein the step of dividing an operation into a set
2 of work partitions includes generating a plurality of query fragments for at least a portion of said
3 operation.

1 26. (Unchanged) The method of Claim 25 wherein:
2 the operation corresponds to at least a portion of a query; and
3 the plurality of query fragments are generated based on said query.

1 27. (Once Amended) The method of Claim 20 wherein the step of assigning work partitions
2 is performed by assigning the work partitions in a sequence based at least in part on sizes
3 associated with the work partitions.

1 28. (Unchanged) The method of Claim 27 wherein the step of assigning the work partitions
2 in a sequence is performed by assigning relatively larger work partitions before assigning
3 relatively smaller work partitions.

1 29. (Once Amended) The method of Claim 22 wherein:

2 the operation is specified in a query that corresponds to a hierarchy of operations; and
3 the step of assigning said at least one entity a second work partition includes
4 determining whether there are any unassigned work partitions from a first level in
5 the hierarchy to which said first work partition belonged; and
6 if there are no unassigned work partitions from the first level in the hierarchy,
7 then selecting said second work partition from a level in said hierarchy
8 that is two levels belowabove said first level in said hierarchy.

1 30. (Unchanged) The method of Claim 20 wherein:
2 the method includes the step of generating a serial execution plan for operations in a
3 database management system (DBMS) running on a computer system;
4 the method includes the step of generating a parallelized execution plan for said serial
5 execution plan, said parallelized execution plan including first and second
6 operations;
7 the step of dividing an operation is performed by dividing said second operation;
8 the plurality of entities includes one or more slave processes operating on a plurality of
9 data partitions, the quantity of said data partitions being greater than the quantity
10 of said slave processes;
11 executing said parallelized execution plan when a plurality of parallel resources of said
12 computer system are available; and
13 executing said serial execution plan when said plurality of resources are not available.

1 31. (Unchanged) The method of claim 30 wherein said step of generating a parallelized
2 execution plan includes the steps of:

3 identifying one or more segments of said serial execution plan that can be parallelized;

4 and

5 identifying partitioning requirements of said one or more segments

1 32. (Unchanged) The method of claim 30 wherein said step of generating a parallelized
2 execution plan is based on a specification of parallelism in a statement specifying one of said
3 operations.

1 33. (Unchanged) The method of Claim 20 further comprising the steps of:

2 generating an execution plan for said operation;

3 examining said execution plan from bottom up;

4 identifying a parallelized portion of said execution plan, said parallelized portion can be
5 processed in parallel, said parallelized portion including first and second
6 operations, said first and second operations being executable in parallel;
7 wherein the step of dividing the operation is performed by dividing said second
8 operation;

9 wherein the plurality of entities includes one or more slave processes operating on a
10 plurality of data partitions, the quantity of said data partitions being greater than
11 the quantity of said slave processes;

12 identifying some serial portion of said execution plan, said serial portion can be
13 processed in serial;

14 allocating a central scheduler between said parallelized portion and said serial portion.

1 34. (Unchanged) The method of Claim 33 further including the steps of:
2 identifying a first data flow requirement for a first portion of said execution plan said first
3 data flow requirement corresponding to a partitioning of a data flow required by
4 said first portion;
5 identifying a second data flow requirement for a second portion of said execution plan
6 said second data flow requirement corresponding by said second portion; and
7 allocating a data flow director between said first portion and said second portion when
8 said first data flow requirement is not compatible with said second data flow
9 requirement said data flow director repartitioning a data flow of said first portion
10 to be compatible with said second data flow requirement

1 35. (Unchanged) The method of Claim 20 further comprising the steps of:
2 generating an execution plan to execute database management system (DBMS)
3 operations in parallel, said execution plan including first and second operations;
4 wherein the step of dividing said operation is performed by dividing said second
5 operation;
6 initiating an operation coordinator in a computer system to coordinate execution of said
7 execution plan;
8 initiating, by said operation coordinator, a first set of slaves operating on a plurality of
9 data partitions to produce data, the quantity of said data partitions being greater
10 than the quantity of said first set of slave processes;

11 initiating, as said plurality of entities, by said operation coordinator, a second set of
12 slaves to consume data; and
13 directing said second set of slaves to produce data and said first set of slaves to consume
14 data when said first set of slaves finishes producing data.

1 36. (Unchanged) The method of claim 35 wherein said execution plan is comprised of
2 operator nodes and said operator nodes are linked together to form execution sets

1 37. (Unchanged) The method of Claim 20 further comprising the steps of:
2 generating an execution plan to execute said operations in parallel, said execution plan
3 including first and second operations;
4 wherein the step of dividing said operation includes dividing said first operation;
5 initiating a data flow scheduler in said computer system to coordinate data flow;
6 initiating, as said plurality of entities, by said data flow scheduler, producer slaves
7 operating on a plurality of data partitions to produce a first data production;
8 initiating, by said data flow scheduler, consumer slaves to consume said first data
9 production;
10 transmitting a ready message to said data flow scheduler when said producer slaves
11 become ready to produce data;
12 transmitting a completion message to said data flow scheduler when said first data
13 production is completed;
14 generating, by said data flow scheduler, in response to said completion message, an
15 identification of a plurality of said consumer slaves that did not receive data in

16 said first data production, said generating step using information derived from
17 said ready message;
18 examining, by said producer slaves, said identification during a subsequent data
19 production; and
20 reducing said subsequent data production such that said subsequent data production does
21 not produce data for said plurality of said consumer slaves

1 38. (Unchanged) A method for processing a query, the method comprising the steps of:
2 receiving a statement that specifies at least (a) an operation and (b) a degree of
3 parallelism to use in performing the operation;
4 dividing the operation into a set of work partitions;
5 performing a determination of how many entities to use to perform said operation based,
6 at least in part, on the degree of parallelism specified in said statement;
7 assigning work partitions from said set of work partitions to a plurality of entities based
8 on said determination; and
9 said plurality of entities operating in parallel on work partitions assigned to them to
10 perform said operation.

1 39. (Unchanged) The method of Claim 38 wherein:
2 the query requires a plurality of operations; and
3 the statement specifies said degree of parallelism for a subset of the plurality of
4 operations required by the query.

1 40. (Unchanged) The method of Claim 39 wherein the degree of parallelism specified by the
2 query indicates that no amount of parallelism is to be used during execution of a particular
3 portion of the query.

1 41. (Unchanged) The method of Claim 38 wherein the degree of parallelism specified by the
2 query indicates a maximum amount of parallelism to use during execution of said operation.

1 42. (Unchanged) A method of processing a query, the method comprising the steps of:
2 dividing an operation required by said query into a set of work partitions by generating a
3 set of query fragments;
4 incorporating hints into at least some of said query fragments, wherein the hint associated
5 with a given query fragment indicates how to perform the work partition
6 associated with said given query fragment;
7 assigning query fragments from said set of query fragments to a plurality of entities; and
8 said plurality of entities operating in parallel on query fragments assigned to them to
9 perform said operation, wherein entities working on a query fragment associated
10 with a hint perform the work partition associated with said query fragment in a
11 manner dictated by said hint.

1 43. (Unchanged) The method of Claim 42 wherein the step of incorporating hints includes
2 incorporating hints that dictate the operation of a table scan.

1 44. (Unchanged) The method of Claim 43 wherein the step of incorporating hints that dictate
2 the operation of a table scan includes incorporating hints that rowed partitioning is to be used
3 during the table scan.

1 45. (Unchanged) The method of Claim 42 wherein the step of incorporating hints includes
2 incorporating hints that specify performance of a full table scan.

1 46. (Unchanged) The method of Claim 42 wherein the step of incorporating hints includes
2 incorporating hints that specify using a particular type of join.

1 47. (Unchanged) The method of Claim 46 wherein the step of incorporating hints that specify
2 using a particular type of join includes incorporating hints that specify using a sort/merge join.

1 48. (Unchanged) The method of Claim 46 wherein the step of incorporating hints that specify
2 using a particular type of join includes incorporating hints that specify using a nested loop join.

1 49. (Unchanged) A method of processing a query, the method comprising the steps of:
2 determining a hierarchy of operations associated with a query;
3 dividing a first operation required by said query into a first set of work partitions;
4 dividing a second operation required by said query into a second set of work partitions,
5 wherein said second operation immediately follows said first operation in said
6 hierarchy;

dividing a third operation required by said query into a third set of work partitions, wherein said third operation immediately follows said second operation in said hierarchy;

assigning work partitions from said first set of work partitions to a first plurality of entities;

said first plurality of entities operating in parallel on work partitions assigned to them from said first set of work partitions to perform said first operation;

assigning work partitions from said second set of work partitions to a second plurality of entities, wherein said second plurality of entities are different entities than said first plurality of entities; and

said second plurality of entities operating in parallel on work partitions assigned to them from said second set of work partitions to perform said second operation;

assigning work partitions from said third set of work partitions to said first plurality of entities; and

said first plurality of entities operating in parallel on work partitions assigned to them from said third set of work partitions to perform said third operation.

1 50. (Unchanged) The method of Claim 49 further comprising performing the following steps
2 when a given entity in said first set of entities finishes performing a work partition from said first
3 set of work partitions:
4 determining whether there are any unassigned work partitions from said first set of work
5 partitions; and

6 if there are no unassigned work partitions from said first set of work partitions, then
7 assigning the given entity a work partition selected from said third set of work
8 partitions; and
9 if there are unassigned work partitions from said first set of work partitions, then
10 assigning the given entity a work partition selected from said first set of work
11 partitions.

1 51. (Unchanged) The method of Claim 49 wherein the hierarchy includes odd levels and
2 even levels, and the method further comprises the steps of assigning work partitions from odd
3 levels to said first plurality of entities and work partitions from even levels to said second
4 plurality of entities.

1 52. (Unchanged) The method of Claim 49 wherein performing work partitions in said first set
2 of work partitions causes said first set of entities produce output consumed by said second
3 plurality of entities, and performing work partitions in said third set of work partitions causes
4 said first set of entities to consume output produced by said second plurality of entities.

1 53. (Unchanged) A method of processing a query, the method comprising the steps of:
2 determining that, to execute said query, output from a plurality of producer operations is
3 to be supplied to a consumer operation;
4 wherein a first set of entities is responsible for executing a first producer operation of said
5 plurality of producer operations;

6 wherein a second set of entities is responsible for executing a second producer operation
7 of said plurality of producer operations;
8 wherein a third set of entities is responsible for executing said consumer operation;
9 during execution of said query, performing the steps of
10 determining whether any entity in said first set of entities produced output for a
11 particular entity in said third set of entities; and
12 if no entity in said first set of entities produced output for said particular entity in
13 said third set of entities, then communicating to at least one entity in said
14 second set of entities that is responsible for supplying output to said
15 particular entity that said at least one entity need not produce output for
16 said particular entity.

1 54. (Unchanged) The method of Claim 53 wherein the step of determining whether any entity
2 in said first set of entities produced output for a particular entity in said third set of entities is
3 performed by
4 the particular entity monitoring whether it received any output from any entity in said
5 first set of entities;
6 the particular entity generating data that indicates whether it received any output from
7 any entity in said first set of entities; and
8 determining whether any entity in said first set of entities produced output for said
9 particular entity based on said data.

1 55. (Unchanged) The method of Claim 53 further comprising the steps of:

2 maintaining a bit vector, wherein each entity in said third set of entities corresponds to a
3 bit in the bit vector;
4 when all entities in said first set of entities that produce output for said particular entity
5 have completed their portion of said first producer operation, setting the bit, in the
6 bit vector, that corresponds to said particular entity if said particular entity
7 received no output from said first producer operation; and
8 wherein the step of determining whether any entity in said first set of entities produced
9 output for said particular entity includes the step of inspecting said bit vector to
10 determine whether any entity in said first set of entities produced output for said
11 particular entity.

1 56. (Unchanged) A computer-readable medium carrying instructions for parallelizing an
2 operation, the instructions including instructions for performing the steps of:
3 dividing the operation into a set of work partitions;
4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions; and
7 said plurality of entities operating in parallel on work partitions assigned to them to
8 perform said operation.

1 57. (Unchanged) The computer-readable medium of Claim 56 wherein:
2 the instructions further include instructions for performing the step of receiving a query
3 that requires the operation; and

4 the step of dividing the operation is performed in response to processing the query.

1 58. (Unchanged) The computer-readable medium of Claim 56 wherein the step of assigning
2 work partitions includes:

3 assigning said at least one entity a first work partition from said set of work partitions;

4 and

5 after said at least one entity has completed operating on said first work partition,

6 assigning said at least one entity a second work partition from said set of work
7 partitions.

1 59. (Unchanged) The computer-readable medium of Claim 56 wherein said plurality of
2 entities are a plurality of processes.

1 60. (Unchanged) The computer-readable medium of Claim 59 wherein said plurality of
2 processes reside within a single database system.

1 61. (Unchanged) The computer-readable medium of Claim 56 wherein the step of dividing an
2 operation into a set of work partitions includes generating a plurality of query fragments for at
3 least a portion of said operation.

1 62. (Unchanged) The computer-readable medium of Claim 61 wherein:
2 the operation corresponds to at least a portion of a query; and
3 the plurality of query fragments are generated based on said query.

1 63. (Unchanged) The computer-readable medium of Claim 56 wherein the step of assigning
2 work partitions is performed by assigning the work partitions in a sequence based at least in part
3 on sizes associated with the work partitions.

1 64. (Unchanged) The computer-readable medium of Claim 63 wherein the step of assigning
2 the work partitions in a sequence is performed by assigning relatively larger work partitions
3 before assigning relatively smaller work partitions.

1 65. (Once Amended) The computer-readable medium of Claim 58 wherein:
2 the operation is specified in a query that corresponds to a hierarchy of operations; and
3 the step of assigning said at least one entity a second work partition includes
4 determining whether there are any unassigned work partitions from a first level in
5 the hierarchy to which said first work partition belonged; and
6 if there are no unassigned work partitions from the first level in the hierarchy,
7 then selecting said second work partition from a level in said hierarchy
8 that is two levels belowabove said first level in said hierarchy.

1 66. (Unchanged) The computer-readable medium of Claim 56 wherein:
2 the instructions include instructions for performing the step of generating a serial
3 execution plan for operations in a database management system (DBMS) running
4 on a computer system;

5 the instructions include instructions for performing the step of generating a parallelized
6 execution plan for said serial execution plan, said parallelized execution plan
7 including first and second operations;
8 the step of dividing an operation is performed by dividing said second operation;
9 the plurality of entities includes one or more slave processes operating on a plurality of
10 data partitions, the quantity of said data partitions being greater than the quantity
11 of said slave processes;
12 the instructions include instructions for performing the step of executing said parallelized
13 execution plan when a plurality of parallel resources of said computer system are
14 available; and
15 the instructions include instructions for performing the step of executing said serial
16 execution plan when said plurality of resources are not available.

1 67. (Unchanged) The computer-readable medium of claim 66 wherein said step of generating
2 a parallelized execution plan includes the steps of:

3 identifying one or more segments of said serial execution plan that can be parallelized;
4 and
5 identifying partitioning requirements of said one or more segments

1 68. (Unchanged) The computer-readable medium of claim 66 wherein said step of generating
2 a parallelized execution plan is based on a specification of parallelism in a statement specifying
3 one of said operations

1 69. (Unchanged) The computer-readable medium of Claim 56 further comprising instructions
2 for performing the steps of:

3 generating an execution plan for said operation;
4 examining said execution plan from bottom up;
5 identifying a parallelized portion of said execution plan, said parallelized portion can be
6 processed in parallel, said parallelized portion including first and second
7 operations, said first and second operations being executable in parallel;
8 wherein the step of dividing the operation is performed by dividing said second
9 operation;

10 wherein the plurality of entities includes one or more slave processes operating on a
11 plurality of data partitions, the quantity of said data partitions being greater than
12 the quantity of said slave processes;

13 identifying some serial portion of said execution plan, said serial portion can be
14 processed in serial;

15 allocating a central scheduler between said parallelized portion and said serial portion

1 70. (Unchanged) The computer-readable medium of Claim 69 further including instructions
2 for performing the steps of:

3 identifying a first data flow requirement for a first portion of said execution plan said first
4 data flow requirement corresponding to a partitioning of a data flow required by
5 said first portion;

6 identifying a second data flow requirement for a second portion of said execution plan
7 said second data flow requirement corresponding by said second portion; and

8 allocating a data flow director between said first portion and said second portion when
9 said first data flow requirement is not compatible with said second data flow
10 requirement said data flow director repartitioning a data flow of said first portion
11 to be compatible with said second data flow requirement

1 71. (Unchanged) The computer-readable medium of Claim 56 further comprising instructions
2 for performing the steps of:

3 generating an execution plan to execute database management system (DBMS)
4 operations in parallel, said execution plan including first and second operations;
5 wherein the step of dividing said operation is performed by dividing said second
6 operation;
7 initiating an operation coordinator in a computer system to coordinate execution of said
8 execution plan;
9 initiating, by said operation coordinator, a first set of slaves operating on a plurality of
10 data partitions to produce data, the quantity of said data partitions being greater
11 than the quantity of said first set of slave processes;
12 initiating, as said plurality of entities, by said operation coordinator, a second set of
13 slaves to consume data; and
14 directing said second set of slaves to produce data and said first set of slaves to consume
15 data when said first set of slaves finishes producing data.

1 72. (Unchanged) The computer-readable medium of claim 71 wherein said execution plan is
2 comprised of operator nodes and said operator nodes are linked together to form execution sets

1 73. (Unchanged) The computer-readable medium of Claim 56 further comprising instructions
2 for performing the steps of:
3 generating an execution plan to execute said operations in parallel, said execution plan
4 including first and second operations;
5 wherein the step of dividing said operation includes dividing said first operation;
6 initiating a data flow scheduler in said computer system to coordinate data flow;
7 initiating, as said plurality of entities, by said data flow scheduler, producer slaves
8 operating on a plurality of data partitions to produce a first data production;
9 initiating, by said data flow scheduler, consumer slaves to consume said first data
10 production;
11 transmitting a ready message to said data flow scheduler when said producer slaves
12 become ready to produce data;
13 transmitting a completion message to said data flow scheduler when said first data
14 production is completed;
15 generating, by said data flow scheduler, in response to said completion message, an
16 identification of a plurality of said consumer slaves that did not receive data in
17 said first data production, said generating step using information derived from
18 said ready message;
19 examining, by said producer slaves, said identification during a subsequent data
20 production; and
21 reducing said subsequent data production such that said subsequent data production does
22 not produce data for said plurality of said consumer slaves

1 74. (Unchanged) A computer-readable medium storing instructions for processing a query,
2 the instructions including instructions for performing the steps of:
3 receiving a statement that specifies at least (a) an operation and (b) a degree of
4 parallelism to use in performing the operation;
5 dividing the operation into a set of work partitions;
6 performing a determination of how many entities to use to perform said operation based,
7 at least in part, on the degree of parallelism specified in said statement;
8 assigning work partitions from said set of work partitions to a plurality of entities based
9 on said determination; and
10 said plurality of entities operating in parallel on work partitions assigned to them to
11 perform said operation.

1 75. (Unchanged) The computer-readable medium of Claim 74 wherein:
2 the query requires a plurality of operations; and
3 the statement specifies said degree of parallelism for a subset of the plurality of
4 operations required by the query.

1 76. (Unchanged) The computer-readable medium of Claim 75 wherein the degree of
2 parallelism specified by the query indicates that no amount of parallelism is to be used during
3 execution of a particular portion of the query.

1 77. (Unchanged) The computer-readable medium of Claim 74 wherein the degree of
2 parallelism specified by the query indicates a maximum amount of parallelism to use during
3 execution of said operation.

1 78. (Unchanged) A computer-readable medium carrying instructions for processing a query,
2 the instructions including instructions for performing the steps of:
3 dividing an operation required by said query into a set of work partitions by generating a
4 set of query fragments;
5 incorporating hints into at least some of said query fragments, wherein the hint associated
6 with a given query fragment indicates how to perform the work partition
7 associated with said given query fragment;
8 assigning query fragments from said set of query fragments to a plurality of entities; and
9 said plurality of entities operating in parallel on query fragments assigned to them to
10 perform said operation, wherein entities working on a query fragment associated
11 with a hint perform the work partition associated with said query fragment in a
12 manner dictated by said hint.

1 79. (Unchanged) The computer-readable medium of Claim 78 wherein the step of
2 incorporating hints includes incorporating hints that dictate the operation of a table scan.

1 80. (Unchanged) The computer-readable medium of Claim 79 wherein the step of
2 incorporating hints that dictate the operation of a table scan includes incorporating hints that
3 rowed partitioning is to be used during the table scan.

1 81. (Unchanged) The computer-readable medium of Claim 78 wherein the step of
2 incorporating hints includes incorporating hints that specify performance of a full table scan.

1 82. (Unchanged) The computer-readable medium of Claim 78 wherein the step of
2 incorporating hints includes incorporating hints that specify using a particular type of join.

1 83. (Unchanged) The computer-readable medium of Claim 82 wherein the step of
2 incorporating hints that specify using a particular type of join includes incorporating hints that
3 specify using a sort/merge join.

1 84. (Unchanged) The computer-readable medium of Claim 82 wherein the step of
2 incorporating hints that specify using a particular type of join includes incorporating hints that
3 specify using a nested loop join.

1 85. (Unchanged) A computer-readable medium carrying instructions for processing a query,
2 the instructions including instructions for performing the steps of:
3 determining a hierarchy of operations associated with a query;
4 dividing a first operation required by said query into a first set of work partitions;
5 dividing a second operation required by said query into a second set of work partitions,
6 wherein said second operation immediately follows said first operation in said
7 hierarchy;

8 dividing a third operation required by said query into a third set of work partitions,
9 wherein said third operation immediately follows said second operation in said
10 hierarchy;
11 assigning work partitions from said first set of work partitions to a first plurality of
12 entities;
13 said first plurality of entities operating in parallel on work partitions assigned to them
14 from said first set of work partitions to perform said first operation;
15 assigning work partitions from said second set of work partitions to a second plurality of
16 entities, wherein said second plurality of entities are different entities than said
17 first plurality of entities; and
18 said second plurality of entities operating in parallel on work partitions assigned to them
19 from said second set of work partitions to perform said second operation;
20 assigning work partitions from said third set of work partitions to said first plurality of
21 entities; and
22 said first plurality of entities operating in parallel on work partitions assigned to them
23 from said third set of work partitions to perform said third operation.

1 86. (Unchanged) The computer-readable medium of Claim 85 further comprising instructions
2 for performing the following steps when a given entity in said first set of entities finishes
3 performing a work partition from said first set of work partitions:
4 determining whether there are any unassigned work partitions from said first set of work
5 partitions; and

6 if there are no unassigned work partitions from said first set of work partitions, then
7 assigning the given entity a work partition selected from said third set of work
8 partitions; and
9 if there are unassigned work partitions from said first set of work partitions, then
10 assigning the given entity a work partition selected from said first set of work
11 partitions.

1 87. (Unchanged) The computer-readable medium of Claim 85 wherein the hierarchy includes
2 odd levels and even levels, and the instructions further include instructions for performing the
3 steps of assigning work partitions from odd levels to said first plurality of entities and work
4 partitions from even levels to said second plurality of entities.

1 88. (Unchanged) The computer-readable medium of Claim 85 wherein performing work
2 partitions in said first set of work partitions causes said first set of entities produce output
3 consumed by said second plurality of entities, and performing work partitions in said third set of
4 work partitions causes said first set of entities to consume output produced by said second
5 plurality of entities.

1 89. (Unchanged) A computer-readable medium carrying instructions for processing a query,
2 the instructions including instructions for performing the steps of:
3 determining that, to execute said query, output from a plurality of producer operations is
4 to be supplied to a consumer operation;

5 wherein a first set of entities is responsible for executing a first producer operation of said
6 plurality of producer operations;
7 wherein a second set of entities is responsible for executing a second producer operation
8 of said plurality of producer operations;
9 wherein a third set of entities is responsible for executing said consumer operation;
10 during execution of said query, performing the steps of
11 determining whether any entity in said first set of entities produced output for a
12 particular entity in said third set of entities; and
13 if no entity in said first set of entities produced output for said particular entity in
14 said third set of entities, then communicating to at least one entity in said
15 second set of entities that is responsible for supplying output to said
16 particular entity that said at least one entity need not produce output for
17 said particular entity.

1 90. (unchanged) The computer-readable medium of Claim 89 wherein the step of
2 determining whether any entity in said first set of entities produced output for a particular entity
3 in said third set of entities is performed by
4 the particular entity monitoring whether it received any output from any entity in said
5 first set of entities;
6 the particular entity generating data that indicates whether it received any output from
7 any entity in said first set of entities; and
8 determining whether any entity in said first set of entities produced output for said
9 particular entity based on said data.

1 91. (Unchanged) The computer-readable medium of Claim 89 further including instructions
2 for performing the steps of:

3 maintaining a bit vector, wherein each entity in said third set of entities corresponds to a

4 bit in the bit vector;

5 when all entities in said first set of entities that produce output for said particular entity

6 have completed their portion of said first producer operation, setting the bit, in the

7 bit vector, that corresponds to said particular entity if said particular entity

8 received no output from said first producer operation; and

9 wherein the step of determining whether any entity in said first set of entities produced

10 output for said particular entity includes the step of inspecting said bit vector to

11 determine whether any entity in said first set of entities produced output for said

12 particular entity.